

Numerical determination of partial spectrum of Hermitian matrices using a Lánczos method with selective reorthogonalization

Chris Johnson¹ and A. D. Kennedy²

SUPA, NAIS, and EPCC,
 Department of Physics and Astronomy,
 The University of Edinburgh, The King's Buildings,
 Edinburgh, EH9 3JZ, Scotland

Abstract

We introduce a new algorithm for finding the eigenvalues and eigenvectors of Hermitian matrices within a specified region, based upon the LANSO algorithm of Parlett and Scott. It uses selective reorthogonalization to avoid the duplication of eigenpairs in finite-precision arithmetic, but uses a new bound to decide when such reorthogonalization is required, and only reorthogonalizes with respect to eigenpairs within the region of interest. We investigate its performance for the Hermitian Wilson–Dirac operator $\gamma_5 D$ in lattice quantum chromodynamics, and compare it with previous methods.

Keywords: Spectrum, Lanczos, Krylov, Eigenvalue, Eigenvector, Hermitian, LANSO, Lattice.

1. Introduction

1.1. Motivation

The problem of computing part of the spectrum of a large Hermitian matrix is common to many areas of computational science, but the particular application that motivated this work is the computation of the Neuberger operator for lattice QCD (Quantum Chromodynamics being the quantum field theory

¹chrisj@epcc.ed.ac.uk

²adk@ph.ed.ac.uk

of the strong nuclear force). This requires us to evaluate the *signum* function of the “Hermitian Dirac operator” $\gamma_5 D$ corresponding to some discrete lattice Dirac operator D , which is defined by diagonalizing this matrix and taking the *signum* (± 1) of each of its eigenvalues. It is far too expensive to carry out the full diagonalization, so we use a Zolotarev rational approximation for the *signum* function as this can be evaluated just using matrix addition, multiplication, and inversion by using a multi-shift solver for its stable partial fraction expansion [1]. The approximation is expensive for eigenvalues of $\gamma_5 D$ that are very close to zero, and as there are only a relatively small number of these we want to deflate them and take their sign explicitly. For this reason we need to compute the part of the spectrum of $\gamma_5 D$ around zero.

1.2. Outline

We begin surveying some basic properties of symmetric matrices in order to introduce the notation used throughout the paper. A pedagogical review of simple eigensolver methods then follows, which leads on to the derivation of the Lánczos method with an explanation of the problems associated with it when using finite-precision floating point arithmetic. An overview of the *LANSO* algorithm of Parlett and Scott [2] is introduced which forms the starting point for the work described here. The goal of the algorithm we introduce in this paper is not to find the full spectrum of a large Hermitian matrix, but to find that part of the spectrum lying within some specified range. For the application described in §5 its implementation in Chroma [3] performs significantly better than the state-of-the-art Ritz method [4, 5].

2. Hermitian Matrices and the Power Method

2.1. Basic Properties of Symmetric Matrices

A matrix \mathbf{A} is *Hermitian* (with respect to a sesquilinear inner product) if $\mathbf{A} = \mathbf{A}^\dagger$, which means $(\mathbf{u}, \mathbf{Av}) = (\mathbf{A}^\dagger \mathbf{u}, \mathbf{v}) = (\mathbf{Au}, \mathbf{v}) = (\mathbf{v}, \mathbf{Au})^*$, or equivalently $\mathbf{u}^\dagger \cdot \mathbf{Av} = (\mathbf{A}^\dagger \mathbf{u})^\dagger \cdot \mathbf{v} = (\mathbf{Au})^\dagger \cdot \mathbf{v} = (\mathbf{v}^\dagger \cdot \mathbf{Au})^*$. An eigenvalue λ of \mathbf{A} satisfies $\mathbf{Az} = \lambda \mathbf{z}$ where $\mathbf{z} \neq 0$ is the corresponding eigenvector. The eigenvalues are real and the eigenvectors are orthogonal. Any matrix can be reduced to triangular form \mathbf{T} by a unitary (orthogonal) transformation³ (change of

³This is *Schur normal form*, which follows from the Cayley–Hamilton theorem that every matrix satisfies its characteristic equation, and the fundamental theorem of algebra which states that the characteristic polynomial $p(\lambda) = \det(\mathbf{A} - \lambda)$ has exactly $N = \dim(\mathbf{A})$ complex roots, counting multiplicity.

basis), $\mathbf{A} = \mathbf{QTQ}^{-1} = \mathbf{QTQ}^\dagger$. For \mathbf{A} Hermitian $\mathbf{T}^\dagger = (\mathbf{Q}^\dagger \mathbf{A} \mathbf{Q})^\dagger = \mathbf{Q}^\dagger \mathbf{A}^\dagger \mathbf{Q} = \mathbf{Q}^\dagger \mathbf{A} \mathbf{Q} = \mathbf{T}$ it follows that \mathbf{T} is real and diagonal; thus $\mathbf{AQ} = \mathbf{QT}$ so the columns of \mathbf{Q} furnish the orthonormal eigenvectors.

2.2. Power Method

In order to find eigenvalues and eigenvectors numerically one obvious approach is the *Power Method*. An arbitrary starting vector can, in theory, be expanded in the orthonormal eigenvector basis $\{\mathbf{z}_j\}$, $\mathbf{u}_0 = \sum_j \mathbf{z}_j(\mathbf{z}_j, \mathbf{u}_0)$. The matrix \mathbf{A} is applied to \mathbf{u}_0 and the result normalized to get \mathbf{u}_1 , and so forth: $\mathbf{u}_{k+1} = \mathbf{Au}_k / \|\mathbf{Au}_k\|$, where the norm is $\|\mathbf{x}\| = \sqrt{(\mathbf{x}, \mathbf{x})}$. We then find that $\mathbf{u}_k \propto \lambda_1^k \mathbf{z}_1(\mathbf{z}_1, \mathbf{u}_0) + \sum_{j>1} (\lambda_j / \lambda_1)^k \mathbf{z}_j(\mathbf{z}_j, \mathbf{u}_0)$, and as $\lim_{k \rightarrow \infty} (\lambda_j / \lambda_1)^k = 0$ we find $\lim_{k \rightarrow \infty} \mathbf{u}_k = \mathbf{z}_1$ assuming $(\mathbf{z}_1, \mathbf{u}_0) \neq 0$, where we label the eigenpairs such that $|\lambda_1| > |\lambda_2| > \dots > |\lambda_N|$. If the eigenvalue λ_1 is degenerate then \mathbf{u}_k converges to the eigenvector parallel to \mathbf{u}_0 . The rate of convergence is governed by $|\lambda_2 / \lambda_1|^k = e^{-k(\ln|\lambda_1| - \ln|\lambda_2|)}$. If we shift the matrix \mathbf{A} by a constant then we just shift its eigenvalues by the same constant and leave the eigenvectors unchanged; however, such a shift *does* change the rate of convergence of the power method.

3. Krylov Spaces and Lánczos Algorithm

3.1. Krylov Spaces

We consider a sequence of subspaces of increasing dimension n such that the restriction of \mathbf{A} to them converges to \mathbf{A} as $n \rightarrow \infty$. For an $N \times N$ matrix \mathbf{A} , convergence will always occur because the approximations equal \mathbf{A} for $n \geq N$. In many cases of practical interest the matrix approximates some compact linear operator on an ∞ -dimensional Hilbert space, and we expect the convergence to be governed by the properties of the underlying operator.

In practice we usually do not have an explicit matrix representation of the large (sparse) matrix \mathbf{A} , but we merely have some functional “black box” representation that allows us to apply it to a vector in \mathbb{R}^N . Almost the only spaces we can construct from this are the *Krylov spaces* $\mathcal{K}_n(\mathbf{A}, \mathbf{u}) = \text{span}(\mathbf{u}, \mathbf{Au}, \mathbf{A}^2\mathbf{u}, \dots, \mathbf{A}^{n-1}\mathbf{u})$ where \mathbf{u} is some more-or-less arbitrary starting vector. The only simple generalization is *block Krylov spaces* where we start from more than one vector.

3.2. Arnoldi Method

The vectors $\{\mathbf{A}^j \mathbf{u}\}$ do not form an orthonormal basis for the Krylov space. Furthermore, the corresponding unit vectors $\mathbf{A}^j \mathbf{u} / \|\mathbf{A}^j \mathbf{u}\|$ converge to the largest eigenvector of \mathbf{A} , as they are just successive iterates of the power method. They therefore provide a particularly *bad* choice of basis for numerical computations. It is natural to construct a better orthonormal basis by deflation and normalization,

$$\mathbf{q}_1 = \mathbf{u} / \|\mathbf{u}\|, \quad \mathbf{u}_{j+1} = \mathbf{A} \mathbf{q}_j - \sum_{k=1}^j \mathbf{q}_k (\mathbf{q}_k, \mathbf{A} \mathbf{q}_j), \quad \mathbf{q}_{j+1} = \frac{\mathbf{u}_{j+1}}{\|\mathbf{u}_{j+1}\|};$$

in other words the Gram–Schmidt procedure. This is called the *Arnoldi method*. We see immediately that $(\mathbf{q}_{j+1}, \mathbf{A} \mathbf{q}_j) = (\mathbf{q}_{j+1}, \mathbf{u}_{j+1}) = \|\mathbf{u}_{j+1}\|$. The $n \times n$ matrix \mathbf{Q} whose columns are⁴ $\mathbf{Q} \mathbf{e}_j = \mathbf{q}_j$ therefore furnishes an orthogonal projector $\mathbf{Q} \mathbf{Q}^\dagger = \sum_{j=1}^n \mathbf{q}_j \otimes \mathbf{q}_j^\dagger$ onto $\mathcal{K}_n(\mathbf{A}, \mathbf{u})$.

The restriction of \mathbf{A} to the Krylov space is *Hessenberg* by construction:

$$\mathbf{H} = \mathbf{Q}^\dagger \mathbf{A} \mathbf{Q} = \begin{pmatrix} H_{1,1} & H_{1,2} & & H_{1,n-2} & H_{1,n-1} & H_{1,n} \\ H_{2,1} & H_{2,2} & \cdots & H_{2,n-2} & H_{2,n-1} & H_{2,n} \\ 0 & H_{3,2} & & H_{3,n-2} & H_{3,n-1} & H_{3,n} \\ \vdots & & \ddots & & \vdots & \\ 0 & 0 & & H_{n-1,n-2} & H_{n-1,n-1} & H_{n-1,n} \\ 0 & 0 & \cdots & 0 & H_{n,n-1} & H_{n,n} \end{pmatrix}.$$

We can diagonalize this matrix using the QR algorithm [6] to obtain $\Theta = \mathbf{S}^\dagger \mathbf{H} \mathbf{S}$, where Θ is the diagonal matrix of *Ritz values*, $\Theta_{ij} = \theta_j \delta_{ij}$, and \mathbf{S} the $n \times n$ unitary (orthogonal) matrix whose columns are the corresponding *Ritz vectors* $\mathbf{s}_j = \mathbf{S} \mathbf{e}_j$. We may hope that some of the Ritz values approximate the eigenvalues of \mathbf{A} , $\theta_j \approx \lambda_j$, and that some of the Ritz vectors approximate its eigenvectors, $\mathbf{Q} \mathbf{S} \mathbf{e}_j = \mathbf{Q} \mathbf{s}_j = \mathbf{y}_j \approx \mathbf{z}_j$, provided that the *residual* $\mathbf{R} \equiv \mathbf{A} \mathbf{Q} - \mathbf{Q} \mathbf{H}$ is small, since $\mathbf{A}(\mathbf{Q} \mathbf{S}) = (\mathbf{Q} \mathbf{H} + \mathbf{R}) \mathbf{S} = (\mathbf{Q} \mathbf{S}) \Theta + \mathcal{O}(\|\mathbf{R}\|)$.

3.3. Lánczos Algorithm

We are interested the special case of the Arnoldi method for a Hermitian matrix \mathbf{A} , which means that the matrix \mathbf{H} is also Hermitian, $\mathbf{H}^\dagger = (\mathbf{Q}^\dagger \mathbf{A} \mathbf{Q})^\dagger =$

⁴ \mathbf{e}_j is a basis vector whose components are $[\mathbf{e}_j]_i = \delta_{ij}$.

$Q^\dagger A^\dagger Q = H$. A matrix which is both Hessenberg and Hermitian is *tridiagonal*

$$H = Q^\dagger A Q = \begin{pmatrix} \alpha_1 & \beta_1 & 0 & 0 & 0 & 0 \\ \beta_1 & \alpha_2 & \beta_2 & \cdots & 0 & 0 \\ 0 & \beta_2 & \alpha_3 & & 0 & 0 \\ \vdots & & \ddots & & \vdots & \\ 0 & 0 & 0 & & \alpha_{n-2} & \beta_{n-2} \\ 0 & 0 & 0 & \cdots & \beta_{n-2} & \alpha_{n-1} \\ 0 & 0 & 0 & & 0 & \beta_{n-1} \\ & & & & & \alpha_n \end{pmatrix},$$

where $\beta_j = \|\mathbf{u}_{j+1}\| = (\mathbf{q}_{j+1}, A\mathbf{q}_j)$ and $\alpha_i = (\mathbf{q}_i, A\mathbf{q}_i)$ are real.

We thus have a three-term recurrence relation

$$A\mathbf{q}_j = \beta_j \mathbf{q}_{j+1} + \alpha_j \mathbf{q}_j + \beta_{j-1} \mathbf{q}_{j-1}; \quad (1)$$

this defines the *Lánczos algorithm*. This greatly simplifies the computation; not only is it easier to diagonalize a tridiagonal matrix using the QR algorithm, but also means that $A\mathbf{q}_j$ is automatically (implicitly) orthogonal to all \mathbf{q}_i except for \mathbf{q}_{i-1} , \mathbf{q}_i , and \mathbf{q}_{i+1} . Unfortunately, floating-point arithmetic does not respect implicit orthogonality.

3.4. Loss of orthogonality among the Lánczos vectors

As noted in the previous section, with a basic implementation of the Lánczos algorithm, orthogonality amongst the Lánczos vectors is lost due to rounding errors. The most obvious indication of this loss of orthogonality is the appearance of spurious copies of eigenvectors. It is interesting to store the Lánczos vectors to measure the loss of orthogonality directly, as it allows us to see where the loss of orthogonality occurs. The results are as expected: with the basic Lánczos algorithm (*i.e.*, one with no reorthogonalization) the orthogonality of a Lánczos vector with respect to those calculated more than two steps previously is only implicit; consequently, as rounding errors inevitably bring back components of the early Lánczos vectors, there is nothing to suppress these components. They therefore grow in an unrestrained manner until eventually orthogonality between the most recent Lánczos vector and those calculated early on in the procedure is completely lost. This was demonstrated in [7] where⁵ $\log_{10}(\mathbf{q}_j^* \mathbf{q}_k / \varepsilon)$ is displayed as a symmetric array of numbers with small values representing mutually orthogonal vectors and large ones representing pairs of vectors with a large overlap.

⁵The “unit of least precision”, ε , is the smallest number such that $1 \oplus \varepsilon \neq 1$ in floating-point arithmetic, it is approximately 10^{-7} for single precision and 10^{-14} for double precision.

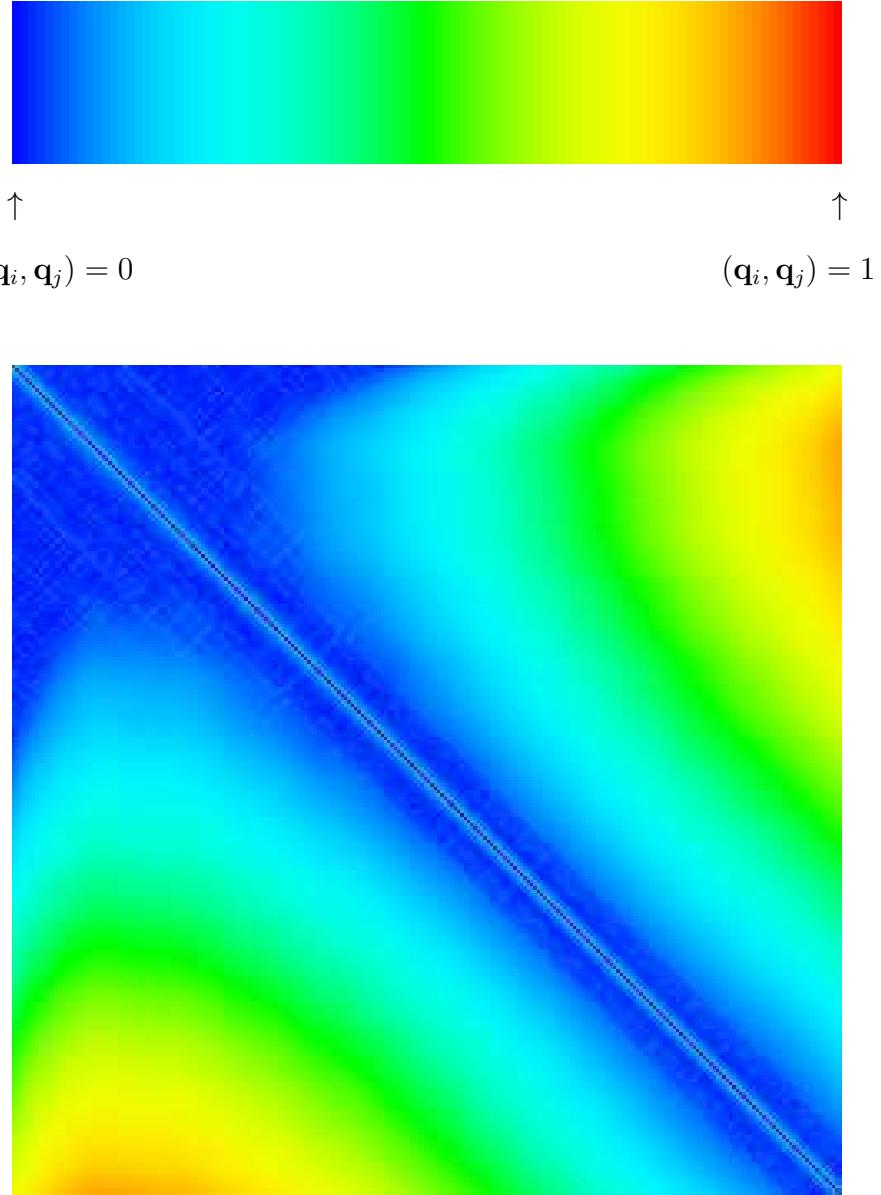


Figure 1: The orthogonality of the Lánczos vectors without reorthogonalization. Increasing Lánczos iterations, i and j are shown in the \downarrow and \rightarrow directions.

For larger systems we can view this as a colour map, an example of which is shown in Figure 1. Large values are represented at the red end of the spectrum and small values at the blue end. Thus the diagonal is shown in red (representing $\|q_j\| = 1$) and mutually orthogonal vectors are shown

in blue.

We can see clearly that with no reorthogonalization, after sufficient steps the new Lánczos vectors lose orthogonality to the early ones. Note that \mathbf{q}_j loses orthogonality with the very earliest Lánczos vectors to a lesser extent compared with those which occur after a few steps. This is to be expected as the initial random starting vector \mathbf{q}_1 will in general not contain large components of any particular eigenvector. However, after a few steps the Lánczos vectors will start to contain large components of the dominant eigenvectors according to the argument given in §2.2 for the power method, and it is precisely these dominant eigenvectors that will grow from rounding errors and so reappear in \mathbf{q}_j .

3.5. Degenerate Eigenspaces and Restarting

In exact arithmetic only one eigenvector will be found for each distinct eigenvalue: if an eigenvalue is degenerate then this vector will be the projection of the initial vector onto its eigenspace. In floating-point arithmetic rounding errors will eventually cause the other eigenvectors to appear; this will take longer in higher-precision arithmetic. This may perhaps be viewed as a case where using floating-point arithmetic is an advantage. Such degenerate eigenvectors can also be found by restarting the Lánczos algorithm with a new initial vector and deflating with respect to the previously known good eigenvectors. This can be repeated until no more degenerate eigenvectors are found. Presumably a block version of the algorithm could be used too, but the choice of block size is not obvious if the maximum degeneracy is not known *a priori*. A cluster of nearby eigenvalues behaves just like a degenerate subspace until sufficient accuracy to resolve the eigenvalues has been attained.

4. Selective Reorthogonalization

We will deem a Ritz vector $\mathbf{y}_j \in \mathcal{K}_n(\mathbf{A}, \mathbf{u})$, where to be “good” if it lies within the Krylov subspace $\mathcal{K}_{n'}(\mathbf{A}, \mathbf{u})$ with $n' < n$, that is if $(\mathbf{y}_j, \mathbf{q}_k) = (\mathbf{Qs}_j, \mathbf{Qe}_k) = (\mathbf{s}_j, \mathbf{e}_k) \approx 0$ for $k > n'$; eigenvalues that are not good will be called “bad”. Paige [8] has shown that the loss of implicit orthogonality occurs primarily in the direction of good Ritz vectors. This is not surprising: if $\mathbf{q}_{n'+1}$ and $\mathbf{q}_{n'+2}$ are orthogonal to an eigenvector \mathbf{z} of \mathbf{A} with eigenvalue λ then all future Lánczos vectors will also be orthogonal to \mathbf{z} in exact arithmetic. We may prove this by induction: assume $(\mathbf{z}, \mathbf{q}_k) = (\mathbf{z}, \mathbf{q}_{k+1}) = 0$ for some $k > n'$, then

$$(\mathbf{z}, \mathbf{Aq}_{k+1}) = (\mathbf{Az}, \mathbf{q}_{k+1}) = \lambda(\mathbf{z}, \mathbf{q}_{k+1}) = 0$$

$$= (\mathbf{z}, \beta_{k+1}\mathbf{q}_{k+2} + \alpha_{k+1}\mathbf{q}_{k+1} + \beta_k\mathbf{q}_k) = \beta_{k+1}(\mathbf{z}, \mathbf{q}_{k+2}),$$

hence $(\mathbf{z}, \mathbf{q}_{k+2}) = 0$ unless the Lánczos process terminates because $\beta_{k+1} = 0$. Concomitantly, any rounding errors that appear in the computation of \mathbf{q}_j for $j > n' + 2$ with a component parallel to \mathbf{z} will not be suppressed by orthogonalization to the previous two Lánczos vectors; moreover, this component will grow as $|\lambda/\lambda'|^k$ where λ' is the largest “bad” eigenvalue.

It therefore suffices to orthogonalize the current Lánczos vectors \mathbf{q}_n and \mathbf{q}_{n+1} explicitly with respect to good eigenvectors sufficiently frequently. This is much cheaper than explicitly orthogonalizing with respect to all the previous Lánczos vectors at each step as in the Arnoldi method.

4.1. LANSO

How often do we need to carry out this reorthogonalization? As rounding errors are of order ε it seems reasonable to choose to do so when the loss of orthogonality has accumulated to be of $\mathcal{O}(\sqrt{\varepsilon})$. We therefore choose to orthogonalize $\mathbf{q}_{n'}$ and $\mathbf{q}_{n'+1}$ with respect to a good Ritz vector \mathbf{y} when $(\mathbf{y}, \mathbf{q}_{n'}) > \sqrt{\varepsilon}$. In their *LANSO* algorithm Parlett and Scott [2] introduce two bounds,

1. The τ bound, $\tau_{ij} \geq |(\mathbf{y}_i, \mathbf{q}_j)|$, that is used to trigger reorthogonalization with respect to \mathbf{y}_i . This bound is computed cheaply by a three-term scalar recurrence.
2. The κ bound, $\kappa \geq \|Q^\dagger Q - 1\|$, that is used to trigger a “pause”, namely a search for new good eigenvectors by running the QR algorithm, followed by a reorthogonalization of the last two Lánczos vectors with respect to all good eigenvectors. This is computed by a more complicated scalar recurrence.

4.1.1. Monitoring the κ and τ bounds

The success of the LANSO method hinges on the ability of κ_j and τ_{kj} to bound well enough the actual values $\|1 - Q^*Q\|$ and $(\mathbf{y}_k, \mathbf{q}_j)$ respectively for any Lánczos step j and all good eigenvectors \mathbf{y}_k calculated thus far. For our relatively small test cases we can store all the Lánczos vectors which make up Q , and all the known good eigenvectors. This enables us to calculate the values of $\|1 - Q^\dagger Q\|$ and $(\mathbf{y}_k, \mathbf{q}_j)$ to compare with these bounds. This information is plotted in Figure 2 where the τ bound is plotted together with the value which it is supposed to bound. The points at which the κ bound triggers a pause are also shown. This figure reveals a number of features. Firstly, $\tau_{kj} > |(\mathbf{y}_k, \mathbf{q}_j)|$ as required. However, the bounds appear

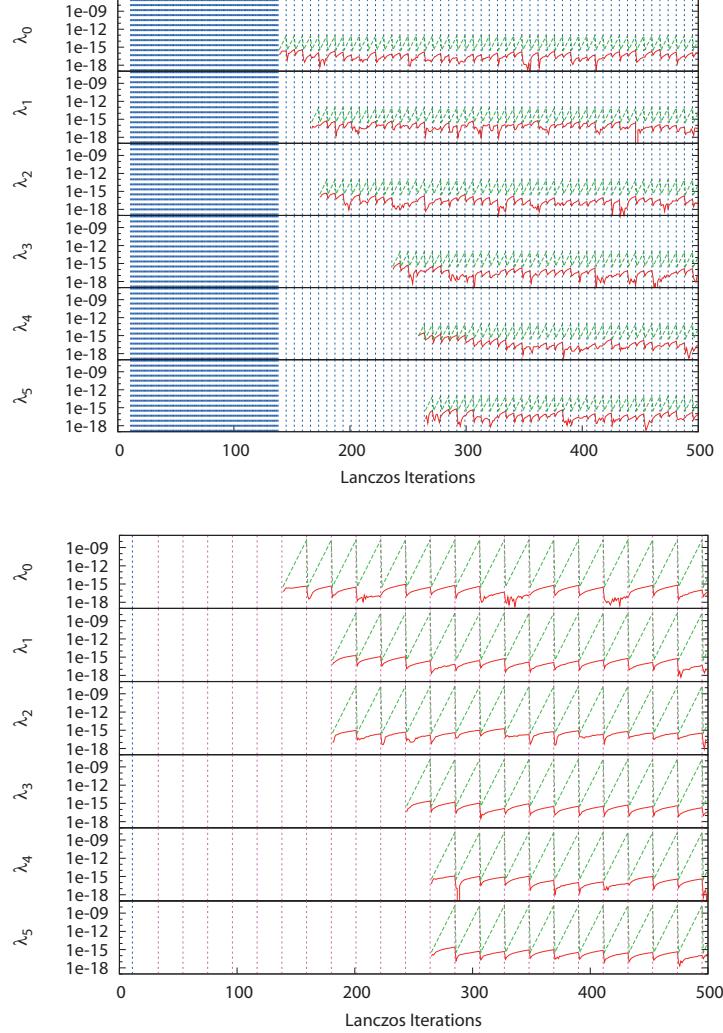


Figure 2: The upper figure plots the values of $|(y_k, q_j)|$ in red and its bound τ_{kj} in green for each y_k for several different Ritz values λ_k . The blue vertical dotted lines show the points at which a diagonalization of H is triggered by the κ bound. The lower figure is similar but here the red vertical lines show the σ bound being used to trigger a full diagonalization. In this example the κ bound was used to trigger the very first pause, but this is not needed: σ could be used from the beginning.

rather pessimistic: the κ bound exceeds the tolerance $\sqrt{\varepsilon}$ and triggers a pause (recall this entails the calculation of the spectrum of the tridiagonal matrix) far more frequently than needed, and the τ bound is often many

orders of magnitude larger than the quantity it is bounding; however, the reorthogonalization triggered by this is relatively inexpensive. Due to the frequent triggering by the κ bound, in practice it is the κ bound and never the individual τ bounds which triggers reorthogonalization.

4.2. New Algorithm

In our application, as in many others, we do not need to find all the eigenpairs: it suffices to find those in a pre-specified region Σ of the spectrum. We only need search for eigenvalues in Σ and selectively reorthogonalize Lánczos vectors with respect to them; we are not interested if duplicate eigenvectors occur outside Σ . In passing, we note that it is easy to restrict the QR iteration to search in the region by a judicious choice of shifts (see §6.2).

Our algorithm replaces both LANSO bounds with a bound σ that is a generalization of the τ bound. σ bounds the loss of orthogonality of a Lánczos vector \mathbf{q}_j with respect to *any* good Ritz vector \mathbf{y} within Σ , even if \mathbf{y} is not explicitly known. We shall require that $\sigma_j \geq \max_{k: \theta_k \in \Sigma} |(\mathbf{y}_k, \mathbf{q}_j)|$, where the maximum is taken over all good Ritz pairs in Σ .

σ is calculated via a three term recurrence relation closely related to that for the τ bound. We consider the propagation and amplification of the lack of orthogonality of the good Ritz vectors with current Lánczos vectors and ignore other inconsequential rounding errors as in [2]. Taking the inner product of (1) with \mathbf{y}_k gives

$$(\mathbf{y}_k, \mathbf{A}\mathbf{q}_j) - (\mathbf{y}_k, \mathbf{q}_{j-1})\beta_{j-1} - (\mathbf{y}_k, \mathbf{q}_j)\alpha_j - (\mathbf{y}_k, \mathbf{q}_{j+1})\beta_j = 0. \quad (2)$$

If $\mathbf{y}_k = \mathbf{Q}\mathbf{s}_k$ is a good Ritz vector within Σ , where (θ_k, \mathbf{s}_k) is a Ritz pair ($\mathbf{H}\mathbf{s}_k = \mathbf{s}_k\theta_k$) then

$$(\mathbf{y}_k, \mathbf{A}\mathbf{q}_j) = (\mathbf{A}\mathbf{y}_k, \mathbf{q}_j) = (\mathbf{A}\mathbf{Q}\mathbf{s}_k, \mathbf{q}_j) = (\mathbf{Q}\mathbf{H}\mathbf{s}_k, \mathbf{q}_j) + (\mathbf{R}\mathbf{s}_k, \mathbf{q}_j) = (\mathbf{y}_k, \mathbf{q}_j)\theta_k, \quad (3)$$

as the residual is orthogonal to the Krylov space, $\mathbf{Q}^\dagger \mathbf{R} = \mathbf{Q}^\dagger \mathbf{A} \mathbf{Q} - \mathbf{Q}^\dagger \mathbf{Q} \mathbf{H} = 0$. From (2) and (3) we obtain

$$(\mathbf{y}_k, \mathbf{q}_{j+1})\beta_j = (\mathbf{y}_k, \mathbf{q}_j)(\theta_k - \alpha_j) - (\mathbf{y}_k, \mathbf{q}_{j-1})\beta_{j-1}.$$

We assume by induction that $\sigma_i \geq |(\mathbf{y}_k, \mathbf{q}_i)| \forall k : \theta_k \in \Sigma, \forall i \leq j$, hence

$$\begin{aligned} |(\mathbf{y}_k, \mathbf{q}_{j+1})| |\beta_j| &\leq |(\mathbf{y}_k, \mathbf{q}_j)| |\theta_k - \alpha_j| + |(\mathbf{y}_k, \mathbf{q}_{j-1})| |\beta_{j-1}| \\ &\leq \max_{\theta \in \Sigma} \sigma_j |\theta - \alpha_j| + \sigma_{j-1} |\beta_{j-1}|; \end{aligned}$$

so we may define

$$\sigma_{j+1} = \frac{\max_{\theta \in \Sigma} |\theta - \alpha_j| \sigma_j + |\beta_{j-1}| \sigma_{j-1}}{|\beta_j|},$$

where the “initial values” $\sigma_{t-1} = \mathcal{O}(\varepsilon)$ and $\sigma_t = \mathcal{O}(\varepsilon)$ correspond to the lack of orthogonality after selectively orthogonalizing a good Ritz vector using a finite precision arithmetic implementation of Gram–Schmidt, t being the last iteration at which the algorithm was paused to search for new Ritz pairs. We shall not give a detailed analysis of this algorithm here but it is very similar for that for LANSO given in [8].

We are interested in applying our algorithm to low density interior regions of the spectrum. The algorithm is surprisingly effective as we find such interior eigenvalues converge rapidly in a manner reminiscent of extremal eigenvalues. The reason why eigenvalues in low-density regions are so well represented in the Krylov space is explained in [9].

4.2.1. Results of the new algorithm

The lower panel of Figure 2 shows the effect of using σ to trigger a pause. We see immediately that when using the σ procedure, diagonalization of H is performed far less frequently than is the case when using the κ procedure.

5. Calculating low eigenvalues of the Fermion matrix

The Lánczos method itself can be used to diagonalize only Hermitian matrices, but the matrices are not required to be positive definite. The Wilson–Dirac fermion matrix D is not Hermitian, but we can exploit the fact that our matrix is “ γ_5 -Hermitian”, $\gamma_5 \mathsf{D} \gamma_5 = \mathsf{D}^\dagger$, where γ_5 is a product of the four Hermitian gamma matrices $\gamma_5 = \gamma_1 \gamma_2 \gamma_3 \gamma_4$ which satisfy the anticommutation relations $\{\gamma_\mu, \gamma_\nu\} = 2\delta_{\mu\nu}$. This allows us to construct the Hermitian matrix $\gamma_5 \mathsf{D}$.

We are interested in the eigenvalues close to gaps in the spectrum, and for $\gamma_5 \mathsf{D}$ there is such a gap around zero. These eigenvalues map to extremal eigenvalues of $\mathsf{D}^\dagger \mathsf{D} = (\gamma_5 \mathsf{D})^2$, but if we use $\mathsf{D}^\dagger \mathsf{D}$ then we have to consider the extra work involved in resolving the sign of the corresponding eigenvalues of $\gamma_5 \mathsf{D}$. This also involves dealing with any mixing which takes place due to the near degeneracy of the approximate eigenvalues, since eigenvalues λ^2 of $\mathsf{D}^\dagger \mathsf{D}$ might be mixtures of eigenvalues of $\gamma_5 \mathsf{D}$ near either $\pm \lambda$.

When using the Lánczos method we see eigenvalues of both large and small magnitude being resolved first giving two regions in the case of $D^\dagger D$ (corresponding to both large and small eigenvalues), and four regions in the case of $\gamma_5 D$ (corresponding to large and small eigenvalues both positive and negative). Figure 3 shows a bar chart of the relative number of small and large converged eigenvalues (regardless of sign) determined at each pause for both $\gamma_5 D$ and $D^\dagger D$.

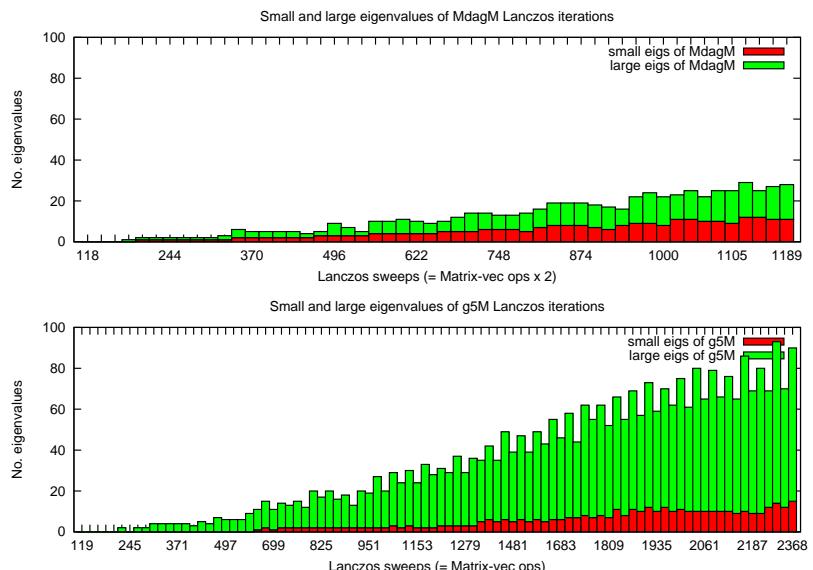


Figure 3: The number of large and small magnitude eigenvalues of $H = \gamma_5 M$ found as a function of the number of Lánczos steps (dimension of the Krylov space). The horizontal axes are scaled to same number of $\gamma_5 D$ applications.

The fact that Figure 3 shows many more large eigenvalues being resolved than small ones gives good motivation for our earlier assertion that we should only look for eigenvalues within the region of interest. If we were to find, construct, and reorthogonalize with respect to all converged eigenvalues at a given Lánczos step most of the time would be spent preserving the orthogonality of regions we are not interested in.

As stated earlier, we are interested in the eigenvalues which are close to a gap in the eigenspectrum around zero. The convergence rates for extremal eigenvalues, *i.e.*, those at either end of the spectrum, are well understood following the work of Kaniel [10], Paige [11] and Saad [12]. This explains why, in the case of $D^\dagger D$ where all the eigenvalues are positive, we see the largest and smallest eigenvalues converge quickly compared with interior ones. In

the case of the matrix $\gamma_5 D$ we see the eigenvalues smallest in magnitude converging quickly. These eigenvalues are not at the extremes of the spectrum but are close to a relatively large void in the spectrum around zero. The convergence rates for such “interior” eigenvalues is explained in [9] where we consider the Kaniel–Paige–Saad bounds applied to the shifted and squared matrix (in this case the optimal shift is zero). Figure 4 shows a comparison of our theoretical bounds with the errors found when finding the eigenvalues close to a gap in the eigenspectrum of the Fermion matrix.

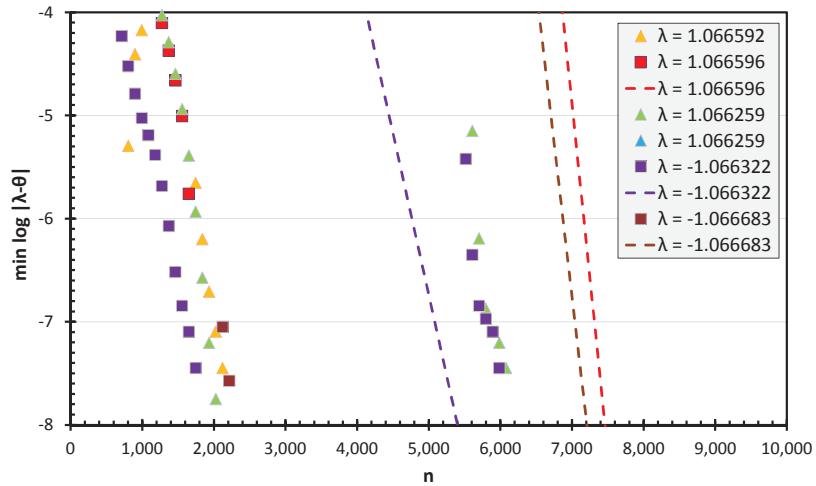


Figure 4: Graph showing the error in eigenvalue estimates as a function of iteration number n (Krylov subspace dimension). We compare this with the theoretical bounds indicated by the dashed lines. The error is determined by taking the absolute value of the difference between the measured Ritz values and the nearest eigenvalue (approximated by the most accurate Ritz value we obtain at the end of the run). The eigenvalue used for a given Ritz value is indicated by different symbols as indicated in the legend. When the error is large this association is somewhat arbitrary but it is unambiguous for the range of errors shown in this graph. The lines correspond to the bounds obtained using the results of [9], again using the spectrum as approximated using the most accurate Ritz values. We see that the purple squares ($\lambda = -1.066322$) and green triangles ($\lambda = 1.066259$) seem to correspond to two orthogonal eigenvectors belonging to degenerate (or very nearly degenerate) eigenvalues. If they were actually degenerate then the second eigenvector would be a fortuitous consequence of rounding error. The agreement between the observed rate of convergence and the theoretical bounds is quite satisfactory.

6. Implementation details

All results obtained here have been obtained using the Chroma [3] package running on 4,096 cores of the UK National Supercomputing Service HECToR [13], after a prototype code was initially implemented in Maple [14]. The Chroma implementation consists of highly optimized parallel linear algebra routines specifically written for lattice QCD, thus we can assume that matrix-vector products, inner-products and general manipulation of vectors and matrices are already optimized. Here we seek to minimize the number of calls to these operations but not to optimize them further. However, we do give some consideration here to patterns of access to large vectors stored in memory, particularly when constructing eigenvectors, and we also consider some optimization of the currently serial diagonalization of the tridiagonal matrix H using the QR method.

6.1. Constructing eigenvectors

Following each application of the QR method, we need to calculate the vectors $\mathbf{y}_i = \mathbf{Q}\mathbf{s}_i$, where \mathbf{s}_i are the columns of \mathbf{S} , *i.e.*, the Ritz vectors. This means that each good Ritz vector \mathbf{y}_i is constructed as a linear combination of Lánczos vectors. The most straightforward method for constructing each eigenvector is via a simple loop as follows

```
DO i = 1 to # good Ritz vectors
  DO j = 1 to # Lanczos vectors
    y[i] = y[i] + q[j] * S[j,i]
  END DO
END DO
```

where the number of good Ritz vectors is expected to be much smaller than the number of Lánczos vectors.

However, this may not be the most efficient ordering. After many Lánczos iterations we will have a large number of Lánczos vectors and they may not all be available in fast memory. We therefore need to ensure that once a Lánczos vector is retrieved from memory we make the most efficient use of it, reducing the need for multiple loads and stores of the vector to and from memory. It may even be that we cannot store all of the Lánczos vectors, and need to reconstruct them on the fly. It therefore makes sense to access (or reconstruct) each Lánczos vector in turn and build up the good Ritz vectors together, by interchanging the order of the loops

```

DO j = 1 to # Lanczos vectors
    Recalculate/access q[j]
    DO i = 1 to # good Ritz vectors
        y[i] = y[i] + q[j] * S[j,i]
    END DO
END DO

```

In both cases the Ritz vectors \mathbf{y}_i are accessed and updated within the inner loop but the second method should result in fewer accesses to the Lánczos vectors, \mathbf{q}_j . Experiments show an average speed-up of approximately 50% in this case.

There are some further interesting architecture-dependent trade-offs that could be investigated. Depending on the amount of memory available and the memory bandwidth we can choose between

1. Storing the Lánczos vectors in main memory (DRAM);
2. Storing the Lánczos vector in secondary storage (disk or Flash RAM);
3. Recomputing the Lánczos vectors at each pause. This minimizes off-chip data transfer, and is “embarrassingly parallel” up to a few global sum operations (for inner products and norms).

A full investigation of these options has not been performed here.

6.2. Diagonalization of \mathbf{H} : QR

We need to pause the Lánczos process periodically to determine the eigen-spectrum of the tridiagonal matrix \mathbf{H} . This can be achieved efficiently using the iterative implicit QR algorithm [6] with suitable shifts.

Many implementations of implicit QR methods exist. The results here were obtained using Lapack [15] routines built on top of BLAS [16], accelerated using the ACML library [17]. The DSTEV Lapack routine could be used to determine all the eigenvalues, and optionally all eigenvectors, of a symmetric tridiagonal matrix. This works well for our needs; however, we are only interested in eigenvalues from within a region Σ , which can give a significant performance benefit. We are better off employing a routine such as DSTEVX which finds eigenvalues only within a specified interval. In the case where Σ is a non-contiguous range, this may result in the routine being called several times, once for each range, or the algorithm could be rewritten to work with a disjoint range. We could also make use of previously known good eigenvalues as shifts, but this has not been implemented.

7. Results

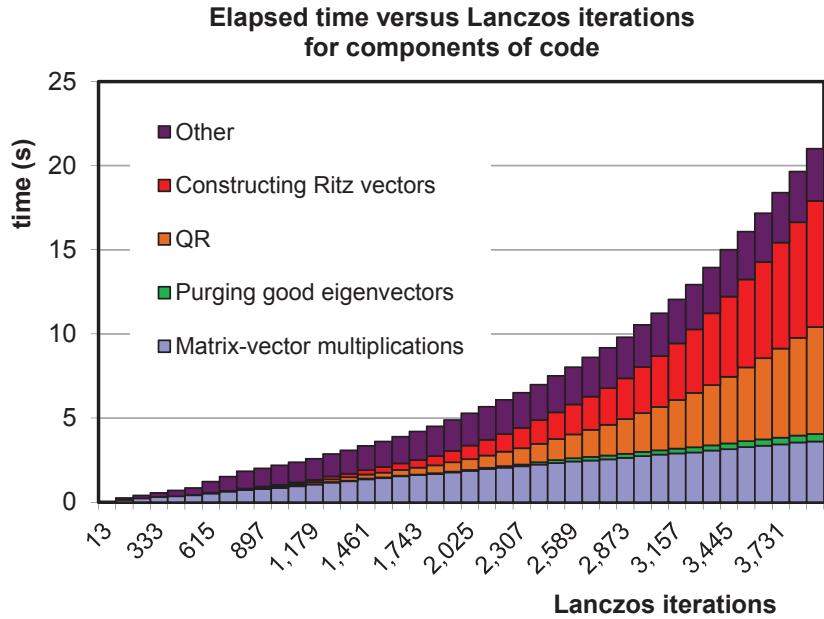


Figure 5: Breakdown of time spent in various parts of the new algorithm versus Lánczos iteration for $\gamma_5 D$ on a $24^3 \times 48$ lattice with 12 degrees of freedom per lattice site (*i.e.*, $N = 7,962,624$) on 4,096 cores of a Cray XT4. “Constructing Ritz vectors” means computing $\mathbf{y} = \mathbf{Q}\mathbf{s}$, and “Purging good eigenvectors” means reorthogonalising the last two Lánczos vectors with all known good Ritz vectors. The x -axis shows the iteration numbers at which the algorithm is paused. The frequency of pauses is such that the x -axis scale is approximately linear.

Figure 5 shows a breakdown of the various components of the algorithm when running on the largest processor count attempted (4,096) for our new algorithm applied to $\gamma_5 D$. We find that with our implementation the most expensive operation is the creation of the eigenvectors of $\gamma_5 D$ followed by the application of the QR method, which is why we wish to create as few eigenvectors as possible. It may also be desirable to implement a faster (e.g., parallel) QR method as the number of eigenvalues required becomes larger.

Figure 6 shows that the speed-up of the creation of eigenvalues with processor count is super-linear. This is due to the fact that with increasing processor count the number of eigenvectors which can be held in cache on each processor

increases as the local sub-vectors become smaller. The net result is a super-linear speed-up of the entire algorithm with processor count.

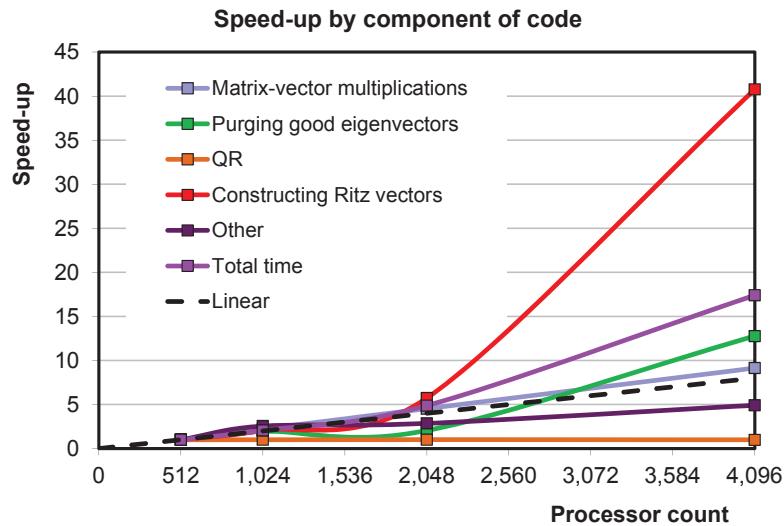


Figure 6: Parallel speed up of new algorithm and its components

In order to illustrate the efficiency of our implementation of our new variant of the LANSO algorithm to find low-lying eigenmodes of the fermion matrix we compare it with the current state-of-the-art, the Chroma implementation of the Kalkreuter–Simma algorithm described in [4]. This method uses a conjugate gradient (CG) method to minimize the Ritz functional $\mu(\mathbf{z}) = (\mathbf{z}, \mathbf{A}\mathbf{z})/\|\mathbf{z}\|^2$ with $\mathbf{A} = (\gamma_5\mathbf{D})^2$, where \mathbf{z} is deflated with respect to all previously computed eigenvectors. The CG minimization alternates with a diagonalization of $\gamma_5\mathbf{D}$ on the subspace of computed eigenvectors to separate eigenvalues of $\gamma_5\mathbf{D}$ of different sign but the same magnitude, taking into account that we may not know the full degenerate subspaces.

Comparing like-with-like for the various methods of determining eigenpairs is not completely straightforward as one has to consider some kind of tolerance within which the eigenvalues are determined. In the case of the Kalkreuter–Simma algorithm convergence is specified by stopping criteria on the CG method, whereas in our new algorithm we determine whether a Ritz pair (θ, \mathbf{y}) has converged by looking at the bottom component of the Ritz vector. Moreover, we continue to refine the eigenpairs at each pause, so their accuracy improves: we could deflate with respect to sufficiently good eigenvectors but we have not studied this option.

We compare the results studying the norm of the residual vector $\|(\mathbf{A} - \theta)\mathbf{y}\|$. We adjust the relevant stopping criteria and tolerances until we see similar magnitudes of this norm and then compare the result in terms of the overall computation time: the results are in Figure 7.

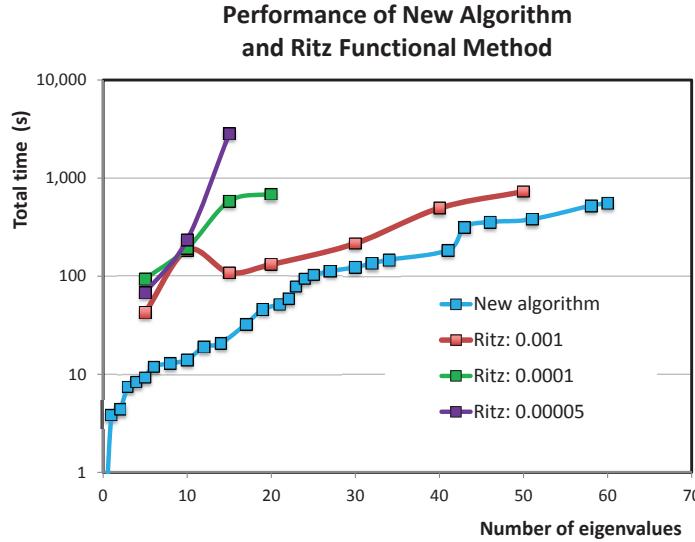


Figure 7: Comparison of our algorithm and the Ritz functional method [4, 5] implemented in Chroma and run on HECToR as a function of the number of small magnitude eigenpairs of $\gamma_5\mathbf{D}$ found. Results are shown for different residual values for the Ritz method; the corresponding errors for our method are always smaller than the best Ritz functional estimates, and decrease as the Krylov space grows.

8. Conclusions

We have introduced a new algorithm to determine the eigenpairs of large Hermitian matrices based on the LANSO method of Parlett and Scott, and implemented and tested it on a realistic large-scale computation in lattice QCD. Our algorithm differs in two ways from LANSO: it only determines eigenpairs within a specified region of the spectrum, as this is all that is needed, and it uses a new σ bound to trigger “pauses” at which Ritz pairs are computed and selective reorthogonalization performed. We found that this reduces the number of such pauses significantly, and moreover far less work is required as we only need to construct the eigenvectors we are interested in.

Our method compares very favourably with the methods that are currently in use, and promises to be useful for other problems such as “low-mode averaging” in QCD calculations as well as in applications in other areas. We have indicated several possible improvements that could be studied in future.

9. Acknowledgements

We would like to thank Balint Joo for his help with Chroma. We gratefully acknowledge the support of the Centre for Numerical Algorithms and Intelligent Software (EPSRC EP/G036136/1) together with STFC (ST/G000522/1) in the preparation of this work.

This work made use of the facilities of HECToR, the UK’s national high-performance computing service, which is provided by UoE HPCx Ltd at the University of Edinburgh, Cray Inc and NAG Ltd, and funded by the Office of Science and Technology through EPSRC’s High End Computing Programme.

References

- [1] A. D. Kennedy, Algorithms for dynamical fermions, World Scientific, 2006, Ch. 2, pp. 15–82.
- [2] B. N. Parlett, D. S. Scott, The Lanczos algorithm with selective orthogonalization, *Mathematics of Computation* 33 (145) (1979) 217–238.
URL <ftp://ftp.math.utah.edu/pub/tex/bib/gvl.bib>; <ftp://ftp.math.utah.edu/pub/tex/bib/mathcomp1970.bib>
- [3] R. G. Edwards, B. Joo, The Chroma software system for lattice QCD, *Nucl. Phys. Proc. Suppl.* 140 (2005) 832.
- [4] T. Kalkreuter, H. Simma, An accelerated conjugate gradient algorithm to compute low lying eigenvalues: A study for the Dirac operator in SU(2) lattice QCD, *Comput. Phys. Commun.* 93 (1996) 33–47.
- [5] B. Bunk, Computing the lowest eigenvalues of the fermion matrix by subspace iterations, *Nucl. Phys. Proc. Suppl.* 53 (1997) 987–989.
- [6] G. H. Golub, C. F. van Loan, *Matrix Computations*, 3rd Edition, Johns Hopkins University Press, 1996.
- [7] H. D. Simon, The Lanczos algorithm with partial reorthogonalization, *Mathematics of Computation* 42 (165) (1984) 115–142.
URL <http://www.jstor.org/stable/2007563>

- [8] B. N. Parlett, The symmetric eigenvalue problem, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1998.
- [9] C. Johnson, A. D. Kennedy, Bounds on the convergence of Ritz values from Krylov subspaces to interior eigenvalues of Hermitian matrices, arXiv math.NA/1110.3034 (2011).
- [10] S. Kaniel, Estimates for some computational techniques in linear algebra, Math. Comp. (1966) 369–378.
- [11] C. C. Paige, The computation of eigenvalues and eigenvectors of very large sparse matrices, Ph.D. thesis, University of London (1971).
- [12] Y. Saad, On the rates of convergence of the Lánczos and the block-Lánczos methods, SIAM J. Numer. Anal. 17 (5) (1980) 687–706.
- [13] HECToR: UK National Supercomputing Service.
URL <http://www.hector.ac.uk/>
- [14] Maple computer algebra system.
URL <http://www.maplesoft.com/>
- [15] E. Anderson, Z. Bai, J. Dongarra, A. Greenbaum, A. McKenney, J. Du Croz, S. Hammerling, J. Demmel, C. Bischof, D. Sorensen, Lapack: a portable linear algebra library for high-performance computers, in: Proceedings of the 1990 ACM/IEEE conference on Supercomputing, Supercomputing '90, IEEE Computer Society Press, Los Alamitos, CA, USA, 1990, pp. 2–11.
URL <http://dl.acm.org/citation.cfm?id=110382.110385>
- [16] Basic Linear Algebra Subprograms.
URL <http://www.netlib.org/blas/>
- [17] AMD core math library.
URL <http://developer.amd.com/libraries/acml/pages/default.aspx>